



2019 INVESTOR MEETING

NAVIN SHENOY

EXECUTIVE VICE PRESIDENT & GENERAL MANAGER
DATA CENTER GROUP

DISCLOSURES

Statements in this presentation that refer to business outlook, future plans and expectations are forward-looking statements that involve a number of risks and uncertainties. Words such as "anticipates," "expects," "intends," "goals," "plans," "believes," "seeks," "estimates," "continues," "may," "will," "would," "should," "could," and variations of such words and similar expressions are intended to identify such forward-looking statements. Statements that refer to or are based on estimates, forecasts, projections, uncertain events or assumptions, including statements relating to total addressable market (TAM) or market opportunity, future products and the expected availability and benefits of such products, and anticipated trends in our businesses or the markets relevant to them, also identify forward-looking statements. Such statements are based on management's expectations as of May 8, 2019, unless an earlier date is indicated, and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in these forward-looking statements. Important factors that could cause actual results to differ materially from the company's expectations are set forth in Intel's earnings release dated April 25, 2019, which is included as an exhibit to Intel's Form 8-K furnished to the SEC on such date. Additional information regarding these and other factors that could affect Intel's results is included in Intel's SEC filings, including the company's most recent reports on Forms 10-K and 10-Q. Copies of Intel's Form 10-K, 10-Q and 8-K reports may be obtained by visiting our Investor Relations website at www.intc.com or the SEC's website at www.sec.gov.

All information in this presentation reflects management's views as of May 8, 2019, unless an earlier date is indicated. Intel does not undertake, and expressly disclaims any duty, to update any statement made in this presentation, whether as a result of new information, new developments or otherwise, except to the extent that disclosure may be required by law.

KEY MESSAGES

THE DATA-CENTRIC OPPORTUNITY IS MASSIVE

LARGEST OPPORTUNITY IN INTEL'S HISTORY, OVER \$200B TAM BY 2023

INDUSTRY MEGA-TRENDS LEVERAGE OUR STRENGTHS

ARTIFICIAL INTELLIGENCE, CLOUD, CLOUDIFICATION OF NETWORK | EDGE

INTEL HAS AN UNPARALLELED ARRAY OF ASSETS TO FUEL GROWTH

PORTFOLIO OF LEADERSHIP PRODUCTS TO MOVE, STORE AND PROCESS DATA



2019 INVESTOR MEETING

INDUSTRY MEGA-TRENDS

GROWTH OF
ARTIFICIAL INTELLIGENCE



PROLIFERATION OF
CLOUD COMPUTING



CLOUDIFICATION OF THE
NETWORK & EDGE



2019 INVESTOR MEETING

EXPLOSION IN DEMAND FOR COMPUTE

INCREASING COMPUTE DEMAND DIVERSIFYING WORKLOAD NEEDS

COMPUTE DEMAND (MIPS)
~60% CAGR

2014 2015 2016 2017 2018 2019 2020 2021 2022 2023

AI

ANALYTICS

HPC

MULTI-CLOUD & ORCHESTRATION

NETWORK

IN-MEMORY DATABASE

VIRTUALIZATION

SECURITY



2019 INVESTOR MEETING

Source: Amalgamation of analyst data and Intel analysis.

LARGEST DATA-CENTRIC OPPORTUNITY IN INTEL HISTORY

DATA-CENTRIC TAM FORECAST
7% CAGR

>\$200B

>\$150B

21%
MSS



GOAL
GROW REVENUE
FASTER THAN TAM




2019 INVESTOR MEETING

Source: 2018 MSS is based on Intel financials. 2023F Si TAM is based on amalgamation of analyst data and Intel analysis, based upon current expectations and available information and are subject to change without notice. Data Center includes Server, Storage, & Network computing, Ethernet/OPA, Silicon Photonics and Memory. Non-Volatile Memory includes NAND and Optane™ technology. IOT includes addressable Logic ASIC/ASSP, MPU, MCU, DSP for Industrial, Transportation, Automated Driving, Retail, Video Surveillance, Healthcare, Public Sector, Office Automation, Gaming and Smart Home. Forecast is soft estimate subject to significant change and revisions.

DATA CENTER GROUP BUSINESS

INTEL DATA CENTER GROUP REVENUE
12% CAGR



- ◇ Cloud SP + Comms SP approaching 70% of DCG revenue
- ◇ 2019 revenue forecast down mid-single digits YOY
 - ◇ Inventory and capacity absorption off of a record 21% growth year
 - ◇ Continued China weakness




2019 INVESTOR MEETING

Source: Intel financials. Forecast is soft estimate subject to significant change and revisions

PUBLIC CLOUD SP GROWTH & DIVERSIFICATION

INTEL PUBLIC CLOUD SP REVENUE
>30% CAGR



INVESTING TO ENABLE NEXT WAVE CSPS
NEXT WAVE GROWTH OF 33% IN 2018

DEEPEN PARTNERSHIPS WITH CSPS
CUSTOM CPUS >55% OF VOLUME IN 2018

PUBLIC CLOUD BUSINESS IS TAM EXPANSIVE
2/3 OF REV IS TAM EXPANSIVE, AND GROWING
(CONSUMER AND NEW CLOUD SERVICES)

PROLIFERATION OF CLOUD COMPUTING ENTERPRISE AND COMMS SERVICE PROVIDERS

DIGITAL TRANSFORMATION CONTINUES

Rakuten

SIEMENS
Healthineers

(SF) EXPRESS

CLOUD SPS INVESTING IN HYBRID CLOUD SOLUTIONS

AWS
OUTPOSTS

aws

DELL TECHNOLOGIES
CLOUD

DELLTechnologies

GOOGLE CLOUD
ANTHOS

 Google Cloud

MICROSOFT AZURE
VMWARE SOLUTIONS

 Microsoft
Azure



2019 INVESTOR MEETING

ARCHITECTING THE DATA-CENTRIC FUTURE

MOVE FASTER

 ETHERNET

 SILICON PHOTONICS


 OMNI-PATH FABRIC

STORE MORE

 OPTANE™ DC 
PERSISTENT MEMORY

 OPTANE™ DC 
SOLID STATE DRIVE

PROCESS EVERYTHING



SOFTWARE & SYSTEM-LEVEL OPTIMIZED




2019 INVESTOR MEETING

APRIL 2ND LAUNCH

DATA-CENTRIC PORTFOLIO

MOVE FASTER

INTEL®
ETHERNET
800 SERIES ADAPTER




STORE MORE

INTEL®
SSDs



INTEL®
OPTANE™ DC
PERSISTENT MEMORY



2ND GENERATION
INTEL®
XEON® SCALABLE



PROCESS EVERYTHING

INTEL®
XEON® D-1600



INTEL®
AGILEX™



“Only one company can introduce technologies across such a broad set of areas – this is unparalleled.”

Mario Morales, IDC

2ND GENERATION INTEL® XEON® SCALABLE PROCESSOR

>50
STANDARD SKUS


DOZENS
CUSTOM SKUS

8 TO 56
CORES PER SOCKET

1 TO 8
SOCKETS


4.5TB
MEMORY PER SOCKET

UP TO 1.33X
AVG. MAINSTREAM
PERF GEN ON GEN


 **HUAWEI**
CLOUD MANAGEMENT
1.42X MORE
VMS
8260+OPTANE PM VS DRAM
CLOUD ORCHESTRATION


NOKIA
VNETWORK GATEWAY
2.0X
5218N+QAT VS 5118
NETWORK

FORTINET
VIRTUAL NG FIREWALL
3.0X
6230N+QAT VS 6230N
SECURITY

 **Tencent Cloud**
CLOUD VIDEO ANALYSIS
3.26X
8260 DLBOOST VS FP32
AI

 **宝信软件
BAOSIGHT**
INDUSTRIAL BIG DATA
1.39X
8280+OPTANE PM VS DRAM
ANALYTICS

 **redis**
IMDB
UP TO **8X** MORE
VM INSTANCES
8280+OPTANE PM VS DRAM
VIRTUALIZATION



 **LS-DYNA**
PHYSICS SIMULATION
2.01X
9242 VS 8160

GBASE®
IMDB
1.35X
8260+OPTANE PM VS DRAM
IN-MEMORY DATABASE

HPC

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.



2019 INVESTOR MEETING

INTEL® OPTANE™ DC PERSISTENT MEMORY

A PLATFORM APPROACH



INTEL® OPTANE™ DC PERSISTENT MEMORY SAM (2023)

\$10B

>50% CAGR ('18-'23)



IN-MEMORY
DATABASE



VMS, CONTAINERS,
APP DENSITY



CONTENT
DELIVERY



REAL-TIME
ANALYTICS



STORAGE
DATA REPLICATION



HIGH PERFORMANCE
COMPUTING

CUSTOMER PROOF-OF-CONCEPT TRACTION SINCE LAUNCH

>100

FORTUNE 500

5

SUPER 7

>30

NEXT WAVE CSPs

>10

COMMS SPs



2019 INVESTOR MEETING

Source: Amalgamation of analyst data and Intel analysis.

ICE LAKE ON TRACK

2014

2015

2016

2017

2018

2019

2020

2021

2022



**INTEL® XEON®
PROCESSOR E5 V3
HASWELL**



**INTEL® XEON®
PROCESSOR E5 V4
BROADWELL**



**INTEL® XEON®
SCALABLE PROCESSOR
SKYLAKE**



**2ND GEN INTEL® XEON®
SCALABLE PROCESSOR
CASCADE LAKE**

**COOPER LAKE
&
ICE LAKE**

ICE LAKE

PRODUCTION SHIPMENTS 1H'20

SAMPLES SHIPPING NOW

POWERED ON AT MULTIPLE CUSTOMERS



2019 INVESTOR MEETING

INCREASING THE PACE OF INNOVATION

2014

2015

2016

2017

2018

2019

2020

2021

2022

INTEL® XEON®
PROCESSOR E5 V3
HASWELL

INTEL® XEON®
PROCESSOR E5 V4
BROADWELL

INTEL® XEON®
SCALABLE PROCESSOR
SKYLAKE

2ND GEN INTEL® XEON®
SCALABLE PROCESSOR
CASCADE LAKE

COOPER LAKE
&
ICE LAKE

SAPPHIRE
RAPIDS

NEXT
GEN

DRIVING LEADERSHIP WORKLOAD PERFORMANCE

5 TO 7
QUARTER CADENCE

MOVING
TO


4 TO 5
QUARTER CADENCE



2019 INVESTOR MEETING

AI OPPORTUNITY

AI DATA CENTER SI TAM
>20% CAGR



'18 INTEL
DATA CENTER
AI REV

>\$1.7B



2019 INVESTOR MEETING

Source: Amalgamation of analyst data and Intel analysis. Intel AI revenue based on Intel financials and analysis for data center processors and FPGAs

DELIVERING AI COMPUTE FROM EDGE TO CLOUD

FROM CPU TO XPU - ONE SIZE DOES NOT FIT ALL



SCALAR



Intel® Xeon®
Scalable Processor Family



VECTOR



Intel®
Discrete Graphics



SPATIAL



Intel®
FPGA



MATRIX




Intel® Nervana™ NNP
Intel® Movidius™ Myriad™
Intel® Mobileye® EyeQ®

ONEAPI UNIFIED DEVELOPER FRAMEWORK



2019 INVESTOR MEETING

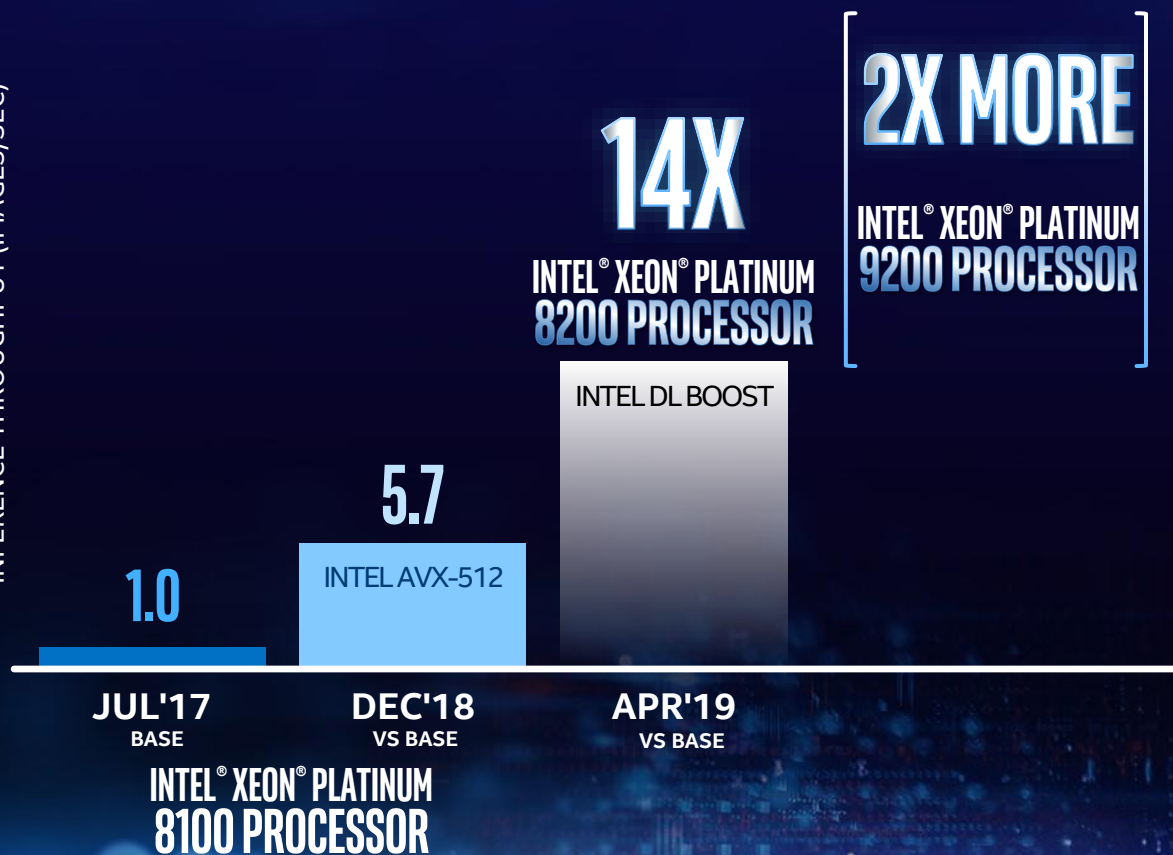


INTEL® DEEP LEARNING BOOST

ONLY CPU WITH BUILT-IN INFERENCE ACCELERATION

INTEL OPTIMIZATION FOR CAFFE RESNET-50

INFERENCE THROUGHPUT (IMAGES/SEC)



Yann LeCun
@ylecun

PyTorch acceleration baked into the latest generation of Intel Xeons. That will help speed up the 200 trillion predictions and 6 billion translations Facebook does every day.
[facebook.com/yann.lecun/pos ...](https://facebook.com/yann.lecun/pos...)

3:18 PM - 9 Apr 2019

SUPPORTED IN ALL MAJOR FRAMEWORKS

Caffe



OpenVINO™

ONNX

PaddlePaddle

PyTorch

TensorFlow



2019 INVESTOR MEETING

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

INTEL® NERVANA™ NEURAL NETWORK PROCESSOR FOR INFERENCE

INCLUDING
ICE LAKE CORES

SILICON
POWERED-ON

PARTNERING WITH
FACEBOOK

INDUSTRY LEADING DEEP LEARNING
TOPS/W & POWER EFFICIENCY

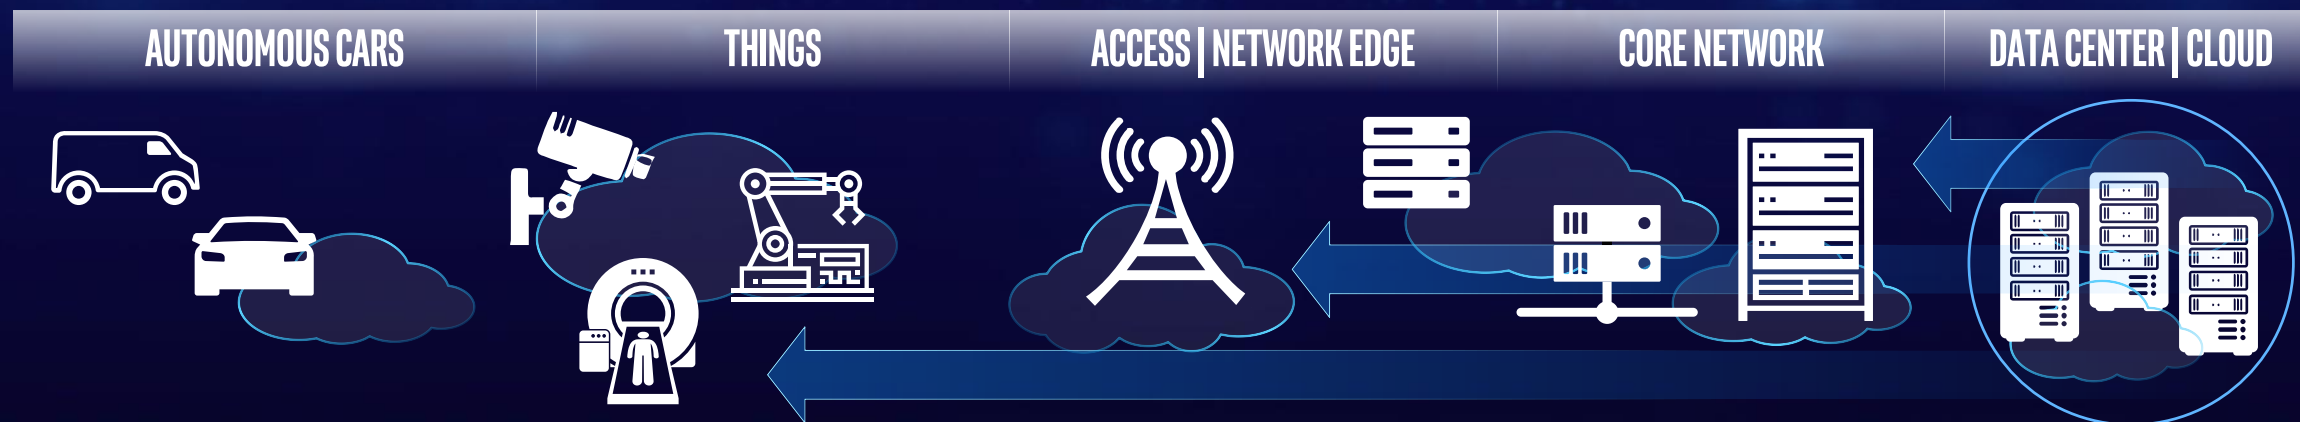


2019 INVESTOR MEETING

Intel® Nervana™ Neural Network Processor for Inference, Tops/W and power efficiency: Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

NETWORK + EDGE COMPUTING

ACCELERATED BY 5G



EDGE COMPUTING		NETWORK EDGE
AUTONOMOUS DRIVING SILICON + SERVICES	IOT SILICON	NETWORK SILICON
~\$10B	~\$30B	~\$25B

~\$65B OPPORTUNITY BY 2023

**2018
INTEL REVENUE
\$9.5B**
>20% GROWTH YOY




2019 INVESTOR MEETING

SANDRA RIVERA

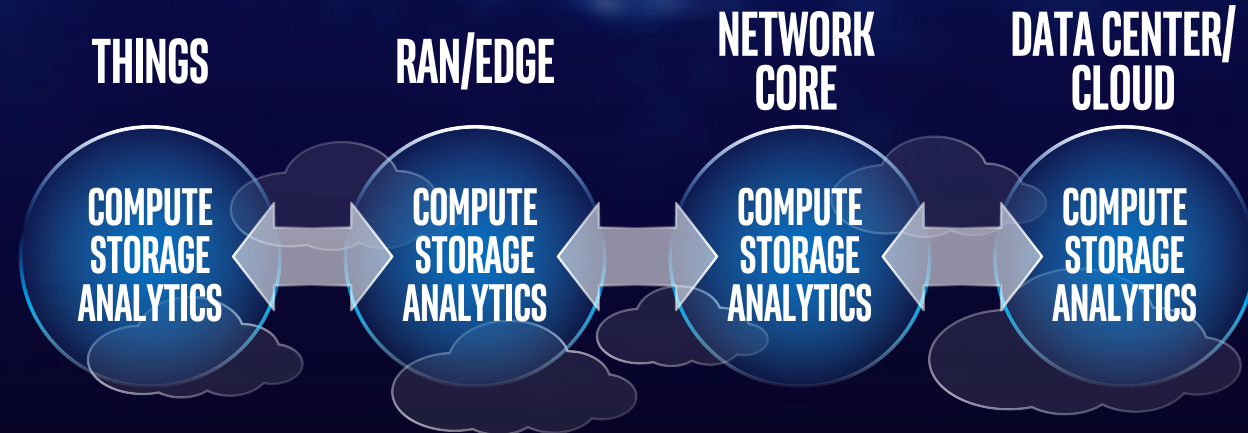
SENIOR VICE PRESIDENT & GENERAL MANAGER
NETWORK PLATFORMS GROUP

CLOUDIFICATION OF THE NETWORK & EDGE

TRADITIONAL NETWORK



INTELLIGENT NETWORK



SCALABILITY & FLEXIBILITY FOR NETWORKING WORKLOADS

intel ETHERNET
intel SILICON PHOTONICS
MOVE

intel OPTANE™ DC
PERSISTENT MEMORY
intel OPTANE™ DC
SOLID STATE DRIVE
STORE

intel XEON™ inside
intel ATOM™ inside
intel AGILEX™ inside
intel CORE™ inside
intel CUSTOM
intel MOVIDIUS™ inside
PROCESS

OpenNESS DPDK
OpenVINO™
SOFTWARE

THE NEXT GENERATION OF NETWORKS IS HERE

2011

NFV
DEFINED

2013

1st NFV
PROOF OF
CONCEPTS

2015

20%
OF COMMS SPS
ADOPT NFV

2017

DPDK
MOVES TO LINUX
FOUNDATION

2018

65%
CORE NW FUNCTIONS
VIRTUALIZED



2019

1st 100%
CLOUD-NATIVE
NETWORK

Rakuten WORLD'S 1ST END TO END CLOUD NATIVE MOBILE NETWORK

100%
ON INTEL ARCHITECTURE

35%
IMPROVED
TCO

4 SKUS
FOR ENTIRE
NETWORK

~1/10TH
OPERATIONS
STAFF

~1 YEAR
FROM CONCEPT TO
DEPLOYMENT

5G READY
NETWORK
ARCHITECTURE

“Our vision is to build a network that innovates at the speed of software and scales at the speed of cloud... leveraging best-in-class technology...to provide a high quality, cost-effective service to our customers.”

Tareq Amin, Group CTIO, Rakuten




2019 INVESTOR MEETING

NETWORK & EDGE GROWTH

ACCELERATED BY 5G

INTEL NETWORK REVENUE
~40% CAGR



AREAS OF FOCUS

DRIVING TRANSFORMATION TO
CLOUD-BASED PLATFORMS

DELIVERING PORTFOLIO OF PRODUCTS FOR
5G AND EDGE

ON TRACK TO BASESTATION MSS
>40% BY 2022




2019 INVESTOR MEETING

Source: Amalgamation of analyst data and Intel analysis. Based on Intel financials for a portion of DCG and PSG.

INTERNET OF THINGS BUSINESS

INTEL IOTG REVENUE
>10% CAGR



AREAS OF FOCUS

AGGREGATION AT THE EDGE

VIDEO INFERENCE

HIGH PERFORMANCE COMPUTE




2019 INVESTOR MEETING

Source: Intel financials. MSS based on amalgamation of analyst data and Intel analysis. ASP & product mix based on 2018 CPU revenue.

INTERNET OF THINGS BUSINESS

INTEL IOTG + AD REVENUE
>15% CAGR



EXTENDING INTO
AUTONOMOUS DRIVING
& DATA SERVICES



2019 INVESTOR MEETING

Source: Intel financials.

KEY MESSAGES

THE DATA-CENTRIC OPPORTUNITY IS MASSIVE

LARGEST OPPORTUNITY IN INTEL'S HISTORY, OVER \$200B TAM BY 2023

INDUSTRY MEGA-TRENDS LEVERAGE OUR STRENGTHS

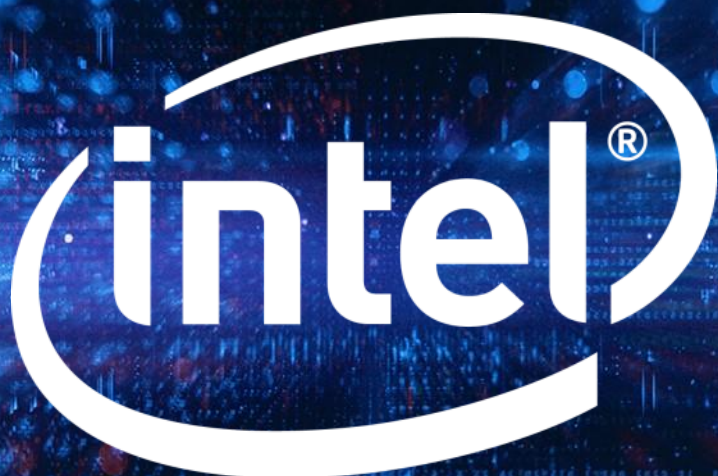
ARTIFICIAL INTELLIGENCE, CLOUD, CLOUDIFICATION OF NETWORK | EDGE

INTEL HAS AN UNPARALLELED ARRAY OF ASSETS TO FUEL GROWTH

PORTFOLIO OF LEADERSHIP PRODUCTS TO MOVE, STORE AND PROCESS DATA



2019 INVESTOR MEETING



CONFIGURATION DISCLOSURE

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Up to 1.33x average generational gains on mainstream Gold SKU: Geomean of est SPECrate2017_int_base, est SPECrate2017_fp_base, Stream Triad, Intel Distribution of Linpack, server side Java. Gold 5218 vs Gold 5118: 1-node, 2x Intel® Xeon® Gold 5218 cpu on Wolf Pass with 384 GB (12 X 32GB 2933 (2666)) total memory, ucode 0x4000013 on RHEL7.6, 3.10.0-957.el7.x86_64, IC18u2, AVX2, HT on all (off Stream, Linpack), Turbo on, result: est int throughput=162, est fp throughput=172, Stream Triad=185, Linpack=1088, server side java=98333, test by Intel on 12/7/2018. 1-node, 2x Intel® Xeon® Gold 5118 cpu on Wolf Pass with 384 GB (12 X 32GB 2666 (2400)) total memory, ucode 0x200004D on RHEL7.6, 3.10.0-957.el7.x86_64, IC18u2, AVX2, HT on all (off Stream, Linpack), Turbo on, result: est int throughput=119, est fp throughput=134, Stream Triad=148.6, Linpack=822, server side java=67434, test by Intel on 11/12/2018.

2.01x LS-Dyna* Explicit, 3car: 1-node, 2x Intel® Xeon® Platinum 8160L cpu on Wolf Pass with 192 GB (12 slots / 16GB / 2666) total memory, ucode 0x200004d on Oracle Linux Server release 7.6 , 3.10.0-862.14.4.el7.crt1.x86_64, Intel SSDSC2BA80, LS-Dyna 9.3-Explicit AVX2 binary, 3car, HT on, Turbo on, test by Intel on 2/26/2019. 1-node, 2x Intel® Xeon® Platinum 9242 cpu on Intel reference platform with 384 GB (24 slots / 16GB / 2933) total memory, ucode 0x4000017 on CentOS 7.6, 3.10.0-957.5.1.el7.x86_64, Intel SSDSC2BA80, LS-Dyna 9.3-Explicit AVX2 binary, 3car, HT on, Turbo on, test by Intel on 3/18/2019.

1.39x BAOSIGHT* xlnsight*: 1-node, 2x Intel® Xeon® Platinum 8260L cpu on S2600WFS with 768 DDR GB (24 slots / 32GB / 2666) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 480GB SSD OS Drive, 1x Intel XC722, xlnsight 2.0 internal workload, HT on, Turbo on, test by Intel/Baosight on 1/8/2019. 1-node, 2x Intel® Xeon® Platinum 8260L cpu on S2600WFS with 192 DDR + 1024 Intel DCPMM GB (12 slots / 16 GB / 2666 DDR + 8 slots / 128 GB / 2666 Intel DCPMM) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 480GB SSD OS Drive, 1x Intel XC722, xlnsight 2.0 internal workload, HT on, Turbo on, test by Intel/Baosight on 1/9/2018.

1.42x Huawei* FusionSphere*: 1-node, 2x Intel® Xeon® Platinum 8260L cpu on Wolf Pass with 1024 GB (16 slots / 64GB / 2666) total memory, ucode 0x400000A on FusionSphere HyperV, 3.10.0-514.44.5.10_96.x86_64 , 1x Intel 800GB SSD OS Drive, 1x Intel 800GB SSD OS Drive, 1x Intel XC722, FusionSphere 6.3.1, mysql-5.7.24, sysbench-1.0.6, HT on, Turbo on, test by Huawei/Intel on 1/11/2018. 1-node, 2x Intel® Xeon® Platinum 8260L cpu on Wolf Pass with 384 DDR + 1536 Intel DCPMM GB (12 slots / 32 GB / 2666 DDR + 12 slots / 128 GB / 2666 Intel DCPMM) total memory, ucode 0x400000A on FusionSphere HyperV, 3.10.0-514.44.5.10_96.x86_64 , 3 x P3520 1.8TB Application Data, 3 x P3520 1.8TB Application Data, 1x Intel XC722, FusionSphere 6.3.1, mysql-5.7.24, sysbench-1.0.6, HT on, Turbo on, test by Huawei/Intel on 1/11/2018.

1.35x GBASE: 1-node, 2x Intel® Xeon® Platinum 8260 cpu on S2600WFT with 768 DDR GB (24 slots / 32GB / 2666) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 400GB SSD OS Drive, 1x Intel XC722, Gbase 8m 6.3.2 OCS Benchmark, HT on, Turbo on, test by GBASE/Intel on 2/19/2019. 1-node, 2x Intel® Xeon® Platinum 8260 cpu on S2600WFT with 192 DDR + 1024 Intel DCPMM GB (12 slots / 16 GB / 2666 DDR + 8 slots / 128 GB / 2666 Intel DCPMM) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 400GB SSD OS Drive, 1x Intel XC722, Gbase 8m 6.3.2 OCS Benchmark, HT on, Turbo on, test by GBASE/Intel on 2/19/2019.

2x Nokia* SDWAN: Configuration #1 (With Intel® QuickAssist® Technology): 2x Intel® Xeon® Gold 5218N Processor on Neon City Platform with 192 GB total memory (12 slots / 16GB / DDR4 2667MHz), ucode 0x4000019, Bios: PLYXCRB 1.86B.0568.D10.1901032132, uCode: 0x4000019 on CentOS 7.5 with Kernel 3.10.0-862, KVM Hypervisor; 1x Intel® QuickAssist Adapter 8970, Cipher: AES-128 SHA-256; Intel® Ethernet Converged Network Adapter X520-SR2; Application: Nokia Nuage SDWAN NSGv 5.3.3U3. Configuration # 2: 2x Intel® Xeon® Gold 5118 Processor on Neon City Platform with 192 GB total memory (12 slots / 16GB / DDR4 2667MHz), ucode 0x4000019, Bios: PLYXCRB 1.86B.0568.D10.1901032132, uCode: 0x4000019 on CentOS 7.5 with Kernel 3.10.0-862, KVM Hypervisor; Intel® Ethernet Converged Network Adapter X520-SR2; Application: Nokia Nuage SDWAN NSGv 5.3.3U3. Results recorded by Intel on 2/14/2018 in collaborate with Nokia.

3.26x latency reduction for Tencent* Cloud Video Analysis: Tested by Tencent as of 1/14/2019. 2 socket Intel® Xeon® Gold Processor, 24 cores HT On Turbo ON Total Memory 192 GB (12 slots/ 16GB/ 2666 MHz), CentOS 7.6 3.10.0-957.el7.x86_64, Compiler: gcc 4.8.5, Deep Learning Framework: Intel® Optimizations for Caffe v1.1.3, Topology: modified inception v3, Tencent's private dataset, BS=1. Comparing performance on same system with FP32 vs INT8 w/ Intel® DL Boost

3x Fortinet* Fortigate*: Configuration #1 (With Intel® QuickAssist Technology) 2x Intel® Xeon® Gold E5-6230N Processor on Neon City Platform with 192 GB total memory (12 slots / 16GB / DDR4 2933MHz), ucode 0x4000019, Bios: PLYXCRB 1.86B.0568.D10.1901032132, uCode: 0x4000019 on CentOS 7.5 with Kernel 3.10.0-862, KVM Hypervisor; 1 x Intel® QuickAssist Adapter 8970, IPsec AES128-SHA256; 1 x Dual Port 40GbE Intel® Ethernet Network Adapter XL710; Application: FortiGate VM64-KVM (v.6.2.0 interim build). Configuration #2 (Without Intel® QuickAssist Technology) : 2x Intel® Xeon® Gold E5-6230N Processor on Neon City Platform with 192 GB total memory (12 slots / 16GB / DDR4 2933MHz), ucode 0x4000019, Bios: PLYXCRB 1.86B.0568.D10.1901032132, uCode: 0x4000019 on CentOS 7.5 with Kernel 3.10.0-862, KVM Hypervisor; 1 x Dual Port 40GbE Intel® Ethernet Network Adapter XL710; Application: FortiGate VM64-KVM (v.6.2.0 interim build). Results recorded by Intel and reviewed by Fortinet on 3/27/2018.

Up to 8X more VMs when running Redis with 8X memory capacity: 1-node, 2x Intel Xeon Platinum 8276 cpu on Intel reference platform with 768 GB (12 slots / 32GB / 2666) total memory, BIOS PLYXCRB1.86B.0573.D10.1901300453 on Fedora-27, 4.20.4-200.fc29.x86_64, 2x40G, Redis 4.0.11, memtier_benchmark-1.2.12 (80/20 read/write); 1K record size, KVM, 1/VM, centos-7.0, HT on, Turbo on, test by Intel on 2/22/2019. 1-node, 2x Intel Xeon Platinum 8276 cpu on Intel reference platform with 192 + 6144 GB (12 slots / 16GB / 2666 DDR + 12 slots / 512GB/ 2666 Intel Optance DCPMM) total memory, BIOS PLYXCRB1.86B.0573.D10.1901300453 on Fedora-27, 4.20.4-200.fc29.x86_64, 2x40G, Redis 4.0.11, memtier_benchmark-1.2.12 (80/20 read/write); 1K record size, KVM, 1/VM, centos-7.0, Memory mode, HT on, Turbo on, test by Intel on 2/22/2019.

CONFIGURATION DISCLOSURE

Intel® Deep Learning Boost

1x inference throughput baseline on Intel® Xeon® Platinum 8180 processor (July 2017) : Tested by Intel as of July 11th 2017: Platform: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). **Performance measured with:** Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50), and https://github.com/soumith/convnet-benchmarks/tree/master/caffe/imagenet_winners (ConvNet benchmarks; files were updated to use newer Caffe prototxt format but are functionally equivalent). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

5.7x inference throughput improvement on Intel® Xeon® Platinum 8180 processor (December 2018) with continued optimizations : Tested by Intel as of November 11th 2018 :2 socket Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz / 28 cores HT ON , Turbo ON Total Memory 376.46GB (12slots / 32 GB / 2666 MHz). CentOS Linux-7.3.1611-Core, kernel: 3.10.0-862.3.3.el7.x86_64, SSD sda RS3WC080 HDD 744.1GB,sdb RS3WC080 HDD 1.5TB,sdc RS3WC080 HDD 5.5TB , Deep Learning Framework Intel® Optimization for caffe version: 551a53d63a6183c233abaa1a19458a25b672ad41 Topology::ResNet_50_v1 BIOS:SE5C620.86B.00.01.0014.070920180847 MKLDNN: 4e333787e0d66a1dca1218e99a891d493dbc8ef1 instances: 2 instances socket:2 (Results on Intel® Xeon® Scalable Processor were measured running multiple instances of the framework. Methodology described here: <https://software.intel.com/en-us/articles/boosting-deep-learning-training-inference-performance-on-xeon-and-xeon-phi>) Synthetic data. Datatype: INT8 Batchsize=64 vs Tested by Intel as of July 11th 2017:2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). **Performance measured with:** Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

14x inference throughput improvement on Intel® Xeon® Platinum 8280 processor with Intel® DL Boost: Tested by Intel as of 2/20/2019. 2 socket Intel® Xeon® Platinum 8280 Processor, 28 cores HT On Turbo ON Total Memory 384 GB (12 slots/ 32GB/ 2933 MHz), BIOS: SE5C620.86B.0D.01.0271.120720180605 (ucode: 0x200004d), Ubuntu 18.04.1 LTS, kernel 4.15.0-45-generic, SSD 1x sda INTEL SSDSC2BA80 SSD 745.2GB, nvme1n1 INTEL SSDPE2KX040T7 SSD 3.7TB, Deep Learning Framework: Intel® Optimization for Caffe version: 1.1.3 (commit hash: 7010334f159da247db3fe3a9d96a3116ca06b09a) , ICC version 18.0.1, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a, model: https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt, BS=64, synthetic Data, 4 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). **Performance measured with:** Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50),. Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

2x More inference throughput improvement on Intel® Xeon® Platinum 9282 processor with Intel® DL Boost : Tested by Intel as of 2/26/2019. Platform: Dragon rock 2 socket Intel® Xeon® Platinum 9282(56 cores per socket), HT ON, turbo ON, Total Memory 768 GB (24 slots/ 32 GB/ 2933 MHz), BIOS:SE5C620.86B.0D.01.0241.112020180249, Centos 7 Kernel 3.10.0-957.5.1.el7.x86_64, Deep Learning Framework: Intel® Optimization for Caffe version: https://github.com/intel/caffe_d554cbf1, ICC 2019.2.187, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a), model: https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt, BS=64, No datalayer syntheticData:3x224x224, 56 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). **Performance measured with:** Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50),. Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

DISCLAIMERS

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks.

Performance results are based on testing as of date specified in the Configuration Disclosure and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

Intel, the Intel logo, EyeQ, Intel Agilex, Intel Atom, Intel Nervana, Intel Optane, Mobileye, Movidius, OpenVINO, Stratix, Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.
© Intel Corporation.